

Sloppy identity

Claire Gardent

Computerlinguistik
Universität des Saarlandes
D-66041 Saarbrücken
E-mail: claire@coli.uni-sb.de

Abstract. Although sloppy interpretation is usually accounted for by theories of ellipsis, it often arises in non-elliptical contexts. In this paper, a theory of sloppy interpretation is provided which captures this fact. The underlying idea is that sloppy interpretation results from a semantic constraint on parallel structures and the theory is shown to predict sloppy readings for deaccented and paycheck sentences as well as relational-, event-, and one-anaphora. It is further shown to capture the interaction of sloppy/strict ambiguity with quantification and binding. Finally, it is compared with other approaches to sloppy identity, in particular [4,12] and [5].

1 Introduction

Sloppy interpretation involves two clauses: a **source** or antecedent clause and a **target** clause that is, a clause containing a proform (the **target proform**). When the antecedent of the target proform also contains a proform (the **source proform**), a sloppy interpretation may arise. In this case, the interpretation of the target proform differs from the interpretation of its source antecedent in the interpretation of the source proform. For instance, in:

Example 1. Jon¹ [washes his₁ car]². Peter does₂ too.

the VP-ellipsis *does* has, among others, the sloppy interpretation *washes Peter's car* (indeed, this is the preferred reading). The interpretation is sloppy because whereas in the source clause *Jon washes his car* the source proform *his* is interpreted as *Jon*, in the target clause *Peter does too*, it is re-interpreted as *Peter*.

Although it is most often associated with VP-ellipsis, the phenomenon of sloppy identity is in fact very pervasive, and can occur in a wide range of configurations for instance: deaccenting (example 2 where the deaccented material is in bold face), paycheck sentences (example 3), VPE as a source proform (examples 4 and 5), non-pronominal referential elements involving implicit arguments (examples 6 and 7), event anaphora (examples 8 and 9) and *one*-anaphora (example 10)¹.

¹ The material occurring between brackets represents the interpretation being considered, in this case, the sloppy interpretation. Most examples are from [15].

Example 2. Jon¹ took his₁ wife to the station. No, BILL **took his wife to the station.** (*Bill took Bill's wife to the station*)

Example 3. Jon¹ spent [his₁ paycheck]² but Peter saved it₂. (*Peter saved Peter's paycheck*)

Example 4. I'll [help you]¹ if you [want me to₁]². I'll kiss you even if you don't₂. (*I'll kiss you even if you don't want me to kiss you*)

Example 5. When Harry [drinks]¹, I always conceal [my belief that he shouldn't₁]². When he gambles, I can't conceal it₂. (*When he gambles, I can't conceal my belief that he shouldn't gamble*)

Example 6. Jon went to a local bar to watch the Super-bowl, and Bob did too. (*Bob went to a bar local to Bob*)

Example 7. George drove to the nearest hospital, and Fred did too. (*Fred drove to the hospital nearest to Fred*)

Example 8. Jon¹ got shot by his₁ father. That happened to Bob too. (*Bob got shot by Bob's father*)

Example 9. Jon¹ kissed his₁ wife, and Bill followed his example (Dahl, 1972) (*Bill kissed Bill's wife*)

Example 10. Although Jon¹ bought a picture of his₁ son, Bill snapped one himself. (*Bill snapped a picture of Bill's son*)

In short, sloppy interpretation can result from many combinations of source and target proform. Nevertheless most theories of sloppiness are restricted either to VPE or to Paycheck Pronouns thus failing to capture this obvious generalisation.

What are the constraints on sloppy interpretation? We claim that parallelism plays a fundamental role in triggering sloppiness. Specifically, it seems that sloppy interpretation is only possible when the antecedent of the source proform has a parallel counterpart in the target clause (cf. [4]). The following examples illustrate this. In example (11), *Jon* has a parallel counterpart in the target clause and consequently, a sloppy interpretation is possible. This sloppy interpretation is however not available in (12) where *Bill* has no parallel counterpart.

Example 11. The policeman who arrested Bill¹ [forgot to read him₁ his₁ rights]² and so did₂ the policeman who arrested Jon.

Example 12. The policeman who arrested Bill¹ [forgot to read him₁ his₁ rights]² and so did₂ Peter.

Of course, sloppy/strict ambiguity is not systematic in that interaction with other linguistic phenomena may block one or the other of the two possible readings. For instance in (13), interaction with quantification results in ruling out a strict reading. Similarly, in example (14), interaction with binding excludes a strict reading. Finally, example (15) shows that the sloppy/strict ambiguity is sensitive to the Pronoun/PN distinction: the source proform must be a pronoun (and not a proper name) for sloppiness to be possible.

Example 13. Every¹ man believes he₁ is a fool. No, PETER believes he is a fool. (sloppy only)

Example 14. Mary persuaded Jon¹ to shave himself₁. No, Mary persuaded PETER to shave himself. (sloppy only)

Example 15. Jon's mother loves Jon and Bill's mother does too (strict only)

In summary, a theory of sloppiness should be general enough to encompass the various configurations in which sloppy/strict ambiguity may arise while correctly accounting for the parallelism constraint illustrated by examples (11) and (12). Furthermore, it should predict the interaction of sloppy/strict ambiguity with such independent phenomena as quantification, binding theory and syntactic categorisation.

In this paper, I present a theory of sloppiness which adheres to these requirements. The proposed analysis is based on a simple parallelism constraint which is claimed to govern the interpretation of parallel propositions. In essence, this constraint requires that parallel structures share a non-trivial part of their semantics. Crucially, this constraint is stated in terms of equations and solved using Higher-Order Unification. It is the use of this particular mechanism which allows us to make the appropriate linguistic predictions: the multiple solutions generated by HOU mirror the sloppy/strict of natural language.

However, HOU is also known to systematically over-generate in that it yields solutions which although they are mathematically valid, are linguistically incorrect. To overcome this problem, we use a form of HOU developed for guiding inductive theorem proving, namely Higher-Order Coloured Unification (HOCU). The basic idea is that HOCU provides a general framework in which to model the interface between semantic construction and other levels of linguistic information. As we shall see, this yields a linguistically plausible way of avoiding over-generation.

The paper is structured as follows. Section 2 sketches the fundamentals of HOCU, and presents the analysis we propose. Section 3 shows that this analysis predicts the fact that both ellipsis and deaccenting may give rise to sloppy/strict ambiguity, section 4 that it naturally extends to paycheck pronouns, and section

5 that it encompasses those cases of sloppy/strict ambiguity in which the source proform is a VP-ellipsis. In section 6, I review the remaining sloppy/strict configurations and briefly indicate how they can be dealt with. Section 7 focuses on the interaction of sloppiness with quantification, binding and syntax. Section 8 compares the approach with previous proposals and in particular [12], [4] and [5]. Section 9 concludes with pointers to further research.

2 The analysis

Our analysis of sloppy identity falls out of an independent constraint on parallel structures. In [8], I argue that this constraint yields a simple treatment of deaccenting and of its relationship to ellipsis. In this paper, I show that it also provides a general theory of sloppy identity.

The parallelism constraint is a semantic constraint which governs the interpretation of parallel structures. Following [4] (henceforth, DSP), we take the structuring of parallelism as given, that is, we assume that parallel elements are known. The language used for semantic representation is the typed lambda calculus and discourse anaphors such as ellipses and discourse pronouns are represented by free variables of the appropriate type. For instance, a VP-ellipsis is assigned the semantic representation $P_{(e,t)}$ that is, a property variable. The parallelism constraint is as follows.

Given $SSem$ and $TSem$, the semantic representations of the source and target utterances and $SP_1, \dots, SP_n, TP_1, \dots, TP_n$, the semantic representations of the source and target parallel elements, the interpretation of parallel utterances must obey the following equations:

$$\begin{aligned} SSem &= A(SP_1), \dots, (SP_n) \\ TSem &= A(TP_1), \dots, (TP_n) \end{aligned}$$

Crucially, these equations are resolved using Higher-Order Coloured Unification. Given the equation $(M = N)$, the unification problem consists in finding a well-formed coloured substitution of terms for free variables that will make M and N equal in the theory of $\alpha\beta\eta$ -identity. For instance, given the equation $l(j, m) = R(j)$, a possible solution is the substitution which assigns $\lambda x.l(x, m)$ to R , written $\{R \leftarrow \lambda x.l(x, m)\}$. As a result, any free variable occurring in the equations (and in particular any free variable representing a discourse anaphor) is assigned a value. As we shall see, this method yields a uniform account of sloppy/strict ambiguity.

What precisely is Higher-Order Coloured Unification (HOCU) and how does it differ from standard HOU? For a precise description of HOCU, the reader is referred to [14]. Briefly, the important difference is that HOCU operates on a variant of the simply typed λ -calculus where symbol occurrences other than bound variables can be annotated with *colours*. Further, colours can be either constants or variables and restrict the unification process as follows. Colour variables unify both with colour variables and with colour constants. In contrast,

a colour constant can only unify with the same colour constant. Further, variables labelled with colour constants are subject to the following restriction.

For any colour constant c and any c -coloured variable V_c , a well-formed coloured substitution must assign to V_c a c -monochrome term i.e., a term whose symbols are c -coloured.

Why use HOCU? As we have already shown (cf. [9]), the intuition is that colours allows us to create an interface between semantic construction and other sources of linguistic knowledge. In this paper, we shall see that it is particularly useful in avoiding over-generation. We now illustrate this point by a simple example. Consider (16).

Example 16. Jon¹ likes his₁ wife. Peter does too.

In an HOU setting, the parallelism constraint yields the following equations for this discourse (recall that ellipses are represented using free variables; thus here R represents the ellipsis *does too*):

$$\begin{aligned} l(j, \text{wife_of}(j)) &= A(j) \\ R(p) &= A(p) \end{aligned}$$

Resolution of the first equation yields a total of 4 solutions, namely:

$$\begin{aligned} \{A \leftarrow \lambda x.l(x, \text{wife_of}(x))\} \\ \{A \leftarrow \lambda x.l(x, \text{wife_of}(j))\} \\ \{A \leftarrow \lambda x.l(j, \text{wife_of}(x))\} \\ \{A \leftarrow \lambda x.l(j, \text{wife_of}(j))\} \end{aligned}$$

Linguistically, only the first two solutions are valid. The last two are invalid because they would yield an incorrect semantics for the discourse in (16), namely:

Jon likes Jon's wife and Jon likes Peter's wife.
Jon likes Jon's wife and Jon likes Jon's wife.

Intuitively, the problem is that the term occurrence representing the source subject appears in the target representation, and that it may not do so. To capture this intuition, DSP introduce the *Primary Occurrence Restriction* (POR) which in essence aims at ensuring that the term(s) representing the source parallel element(s) do not occur in the solution.

Primary Occurrence Restriction

Given a labeling of occurrences as either primary or secondary, the POR excludes from the set of linguistically valid solutions, any solution which contains a primary occurrence.

Within the HOCU framework, the POR finds a natural encoding: primary occurrences are \mathbf{p} -coloured whilst the free variable representing an ellipsis is \mathbf{s} -coloured. Since a well-formed coloured substitution only assigns to a c -coloured free variable (where c is a colour constant), a c -monochrome term, it follows that the

term assigned to an ellipsis variable (R_s) may not contain any primary occurrence (because primary occurrences are p -coloured and $p \neq s$). Coming back to example (16), the HOCU equations are (here and in what follows, we ignore irrelevant colours; A, B are colour variables):

$$\begin{aligned} l(j_p, \text{wife_of}(j_A)) &= A_B(j_p) \\ R_s(p_p) &= A_B(p_p) \end{aligned}$$

For which, the following substitutions are well-formed coloured substitutions

$$\begin{aligned} \{R_s \leftarrow \lambda x.l(x, \text{wife_of}(x)), A_B \leftarrow \lambda x.l(x, \text{wife_of}(x))\} \\ \{R_s \leftarrow \lambda x.l(x, \text{wife_of}(j_s)), A_B \leftarrow \lambda x.l(x, \text{wife_of}(j_A))\} \end{aligned}$$

but not these:

$$\begin{aligned} \{R_s \leftarrow \lambda x.l(j_p, \text{wife_of}(x)), A_B \leftarrow \lambda x.l(j_p, \text{wife_of}(x))\} \\ \{R_s \leftarrow \lambda x.l(j_p, \text{wife_of}(j_s)), A_B \leftarrow \lambda x.l(j_p, \text{wife_of}(j_A))\} \end{aligned}$$

To summarise: because HOU can yield several solutions (rather than a single one), it allows us to capture the sloppy/strict ambiguity displayed in natural language discourse. And because we use HOCU (rather than straight HOU), we can eliminate from the set of solutions, those solutions that are linguistically invalid. For more details on the applications of HOCU to natural language semantics, we refer the reader to ([9]). In what follows, we will omit colours unless they are used for something else than the POR. Concretely, this means that colours will only re-appear in section 7, where we concentrate on the interaction between semantic construction and other sources of linguistic information.

3 Sloppy identity in ellipsis and deaccenting contexts

It has often been observed (cf. [19,1,18]) that VP-ellipsis and deaccenting (i.e. prosodic reduction) share a number of interpretive similarities and in particular that they both give rise to sloppy/strict ambiguity. Thus in (18), the target VP *took his wife to the station* is deaccented, and like the elided VP in (17), it can be interpreted either strictly or sloppily (upper letters indicate prosodic prominence and bold face deaccented material).

Example 17. Jon¹ took his₁ wife to the station. No, BILL did.

Example 18. Jon¹ took his₁ wife to the station. No, BILL **took his wife to the station.**

Interestingly, this is exactly what the parallelism constraint predicts. Thus in the VPE case, the parallelism constraint requires that the following equations hold:

$$\begin{aligned} tk(j, \text{wife_of}(j), s) &= A(j) \\ R(b) &= A(b) \end{aligned}$$

Resolving the first equation yields two possible values for A :

$$\begin{aligned} &\{A \leftarrow \lambda x.tk(x, \text{wife_of}(x), s)\} \\ &\{A \leftarrow \lambda x.tk(x, \text{wife_of}(j), s)\} \end{aligned}$$

And consequently, the target clause receives two interpretations, either $tk(b, \text{wife_of}(b), s)$ or $tk(b, \text{wife_of}(j), s)$ – as required. Now consider the deaccenting case. As for the ellipsis case, we assume that anaphors in the source are resolved whereas discourse anaphors in the target are represented using free variables (alternatively, we could resolve them first and let HOU filter unsuitable resolutions out). Specifically, the target pronoun *his* is represented by the free variable x and the equations to be solved are:

$$\begin{aligned} tk(j, \text{wife_of}(j), s) &= A(j) \\ tk(b, \text{wife_of}(x), s) &= A(b) \end{aligned}$$

The first equation is resolved as before, thereby yielding two possibilities for the second equation:

$$\begin{aligned} tk(b, \text{wife_of}(x), s) &= tk(b, \text{wife_of}(b), s) \\ tk(b, \text{wife_of}(x), s) &= tk(b, \text{wife_of}(j), s) \end{aligned}$$

It follows that x can unify either to j or to b and accordingly, the target pronoun *his* is interpreted as either strict (*Jon*) or sloppy (*Bill*).

Another interesting fact about sloppy interpretation in deaccenting and VP ellipsis contexts is that it may involve an extended domain of licensing that is, a domain that extends beyond the clausal level. For instance, [18] notes that in (19) and (20), both the deaccented *I was bad-mouthing her* and the elliptical *I was* can be assigned the sloppy interpretation: *I was bad-mouthing Sue* even though *Sue* is not part of the anaphoric clause.

Example 19. First, Jon told Mary¹ I was bad-mouthing her₁, and then he told SUE I was.

Example 20. First, Jon told Mary¹ I was bad-mouthing her₁, and then he told SUE I was bad-mouthing her.

More generally, the interpretation of both ellipsis and deaccenting can depend on elements which occur outwith their governing categories. This suggests that the semantic licensing of these two phenomena occurs at the sentential rather than at the clausal level. Again this falls out of the parallelism constraint: since this constraint operates on utterance representations, it is not restricted to clauses but may indifferently apply either to clauses or to sentences. Specifically, the analysis for the above examples goes as follows.

For the ellipsis case, the equations are:

$$\begin{aligned} t(j, m, bm(i, m)) &= A(m) \\ t(x, s, R(i)) &= A(s) \end{aligned}$$

Solving the first equation yields $\lambda z.t(j, z, bm(i, z))$ and $\lambda z.t(j, z, bm(i, m))$ as possible values for A . By substitution and β -reduction $A(s)$ is then either $t(j, s, bm(i, s))$ or $t(j, s, bm(i, m))$ and the resulting equations are resolved by the following substitutions:

$$\begin{aligned} \{x \leftarrow j, R \leftarrow \lambda z.bm(z, m)\} \\ \{x \leftarrow j, R \leftarrow \lambda z.bm(z, s)\} \end{aligned}$$

where the first solution yields the strict reading, the second the sloppy. Resolution of the deaccenting case can similarly be summarised as follows. This time, the equations resulting from the parallelism constraint are:

$$\begin{aligned} t(j, m, bm(i, m)) &= A(m) \\ t(x, s, bm(i, y)) &= A(s) \end{aligned}$$

As before, resolving the first equation yields two possible values for A namely, $\lambda z.t(j, z, bm(i, z))$ and $\lambda z.t(j, z, bm(i, m))$. Similarly, $A(s)$ is then either $t(j, s, bm(i, s))$ or $t(j, s, bm(i, m))$. Consequently, (20) also receives both a strict and a sloppy interpretation.

Finally, deaccenting may involve more complex cases of sloppy interpretation as illustrated by the following examples.

Example 21. Jon¹ said he₁ was clever. No, PETER said he was intelligent.

Example 22. The policeman who arrested Jon¹ forgot to read him₁ his₁ rights and so did the one PETER got collared by.

Example (21) is a straightforward deaccenting example where the deaccented material differs in its lexical realisation from its source counterpart; example (22) is more complex and involves both deaccenting and VP-ellipsis. Both cases trigger a sloppy/strict ambiguity in the interpretation of the target clause.

The important point is that in such cases, ellipsis and deaccenting differ in the semantic relation they require to hold between source and target utterances. In the ellipsis case, the relation is one of syntactic identity: the semantic representation of the ellipsis in the target clause must be identical with the semantic representation of part of the source clause. In the case of deaccenting, the relation is more subtle. Consider (23) for instance.

Example 23. First, JON called Mary a Republican and then PETER insulted her.

[18] argues that in such cases, entailment is involved in licensing prosodic reduction: (23) is licensed by the implication that ‘if x calls y a republican, then x insults y ’ and consequently, the deaccented clause is interpreted as *Peter insulted Mary*. Note that through the entailment, *her* is resolved to *Mary*.

However, Rooth himself notes that the licensing relation between a deaccented sentence and its source, need not always be entailment. Thus in examples (24), (25) and (26), there is clearly no entailment relation holding between source and target clause.

Example 24. He bit her and then SHE punched HIM.

Example 25. Tell me who assaulted whom? HE bit HER.

Example 26. First, a policeman arrested Peter and then PAUL got collared by one.

Rather, the requirement seems to be that the two clauses entail a (reasonably specific) common proposition of a more general form. For instance, the deaccenting in (26) seems to be licensed by the fact that both *x arrest y* and *x collared y* entail that *x did something nasty to y* (note that if I say *First, a policeman arrested Peter and then PAUL was invited to TEA by one*, I have to stress *was invited to TEA*). More specifically, one could follow [19] and argue that deaccenting is licensed provided $(X_S \supset Y) \wedge (X_T \supset Y)$ where X_S, X_T are the semantic representations of the source and target clauses respectively, and Y is such that it subsumes X_S and X_T and it is entailed by X_S and X_T .

Now, whatever the relation is, which holds between a deaccented clause and its source, the above data is clearly problematic for the HOU approach. In such cases, the semantic representation of the deaccented material will differ from that of its parallel counterpart in the source and therefore HOU (which is essentially a matching operation on syntactic structures) will fail thus failing to account for some perfectly acceptable discourses.

There is an intuitively natural solution to this problem. As we already argued in [10,8], a variant of HOU is required to deal with the semantics of deaccenting namely, HOU augmented with logical relations (HOU+R, cf. [16]). Crucially, this form of unification takes into account not only syntactic $\alpha\beta\eta$ -identity, but also any logical relation we care to specify. For instance, if the specified relation is entailment, the equation $(M = N)$ will be solved by HOU+R if either M and N syntactically unify or if it can be proved that M entails N . The point is that whatever the relation is which we find, holds between a deaccented clause and its source, we can integrate it into the HOU framework (provided it's a logical relation). Briefly, the HOU+R approach consists in combining HOU with theorem proving. Specifically, each time an intermediate equation $(M = N)$ of type t is found, the theorem prover is called to try and prove that $R(M, N)$ (where R is the specific relation HOU is incremented with). Thus, if we assume that source and target clause must share a common property, any attempt to solve $(M = N)$ will trigger the attempt to prove that $(M \supset Y) \wedge (N \supset Y)$ where Y subsumes both M and N . For a more detailed presentation of HOU+R, we refer the reader to [16]; for a discussion of how this form of unification can be used for deaccenting, see [10,8]. We now briefly sketch how such cases can be handled within the HOU+R framework. Consider example (22) with sloppy interpretation *The policeman who Peter got collared by forgot to read Peter Peter's rights*. The equations to be resolved are:

$$\begin{aligned}\exists x[po(x) \wedge ar(x, j) \wedge f(x, rd(x, j, j'sr))] &= A(j) \\ \exists x[Q(x) \wedge col(x, p) \wedge R(x)] &= A(p)\end{aligned}$$

Solving the first equation yields (we concentrate on the sloppy reading):

$$\{A \leftarrow \lambda z. \exists x[po(x) \wedge ar(x, z) \wedge f(x, rd(x, z, z'sr))]\}$$

Consequently, the second equation becomes

$$\exists x[Q(x) \wedge col(x, p) \wedge R(x)] = \exists x[po(x) \wedge ar(x, p) \wedge f(x, rd(x, p, p'sr))]$$

For which, a possible solution is: $\{Q \leftarrow \lambda z.po(z), R \leftarrow \lambda z.f(z, rd(z, p, p'sr))\}$
After substitution, the lhs of the equation is then:

$$\exists x[po(x) \wedge col(x, p) \wedge f(x, rd(x, p, p'sr))]$$

That is, the underspecified semantics of the target *so did the one Peter got collared by* has correctly been resolved to *the policeman who arrested Peter forgot to read Peter Peter's rights*. Note however that the equation is not solved. Rather, we've reached a situation in which no free variable occurs but still left and right-hand sides are not identical. Within the HOU+R setting, an attempt will then be made to prove that R holds between both sides of the equation. In this case, the theorem prover must prove that both sides of the equations entail a common more general proposition. Specifically, it must be proved that:

$$\begin{aligned} &\exists x[po(x) \wedge ar(x, p) \wedge f(x, rd(x, p, p'sr)) \supset Y] \wedge \\ &\exists x[po(x) \wedge col(x, p) \wedge f(x, rd(x, p, p'sr)) \supset Y] \end{aligned}$$

Since tableaux are refutation systems, we in fact try to prove the negation of this formula and aim for a contradiction. During the proof process, each conjunctive formula can be broken down into its conjuncts which are then added to the current tableau branch. Conversely, a disjunctive formula triggers a branching of the tree whereby each new branch is labelled with one of the disjuncts. Universal and existential formulas license the addition of an instantiation of this formula to the current branch (whereby the instantiation drops the quantifier and replaces the bound variable by an unused free variable in the case of a universal and by a new skolem term in the case of an existential). Finally, if both X_1 and $\neg X_2$ occur on the same tableau branch, and if a substitution can be found which makes X_1 and X_2 equal, the branch is closed. When all tableau branches are closed, the theorem is proved (for a more precise definition of the tableau method, see e.g. [6]). The tableau for the above example can be sketched as follows. In a first phase, we get (we abbreviate $f(x, rd(x, p, p'sr))$ to $fr(x, p)$):

- (1) $\forall x \forall y[ar(x, y) \supset nasty(x, y)]$
- (2) $\forall x \forall y[col(x, y) \supset nasty(x, y)]$
- (3) $ar(v_1, v_2) \supset nasty(v_1, v_2)$
- (4) $col(v_3, v_4) \supset nasty(v_3, v_4)$
- (5) $\neg(\exists x[po(x) \wedge ar(x, p) \wedge fr(x, p) \supset Y] \wedge \exists x[po(x) \wedge col(x, p) \wedge fr(x, p) \supset Y])$
- (6) $\neg \exists x[po(x) \wedge ar(x, p) \wedge fr(x, p) \supset Y] \mid \neg \exists x[po(x) \wedge col(x, p) \wedge fr(x, p) \supset Y]$

1 and 2 formalise the fact that we assume a context where *if x arrests y, then x is nasty to y* and further *if x collars y, then x is nasty to y*. 3 and 4 are

from 1 and 2 by instantiation (where v_i are new free variables). 5 is the negation of what has to be proved. Since $\neg(X \wedge Y)$ is a disjunctive formula, the tableau branches and its disjuncts $\neg X$ and $\neg Y$ are added to the new branches. We now show the continuing derivation for the left branch (the right branch develops in a similar way). For readability, we add lines 3 and 6 of the beginning tableau to the continuing tableau.

$$\begin{array}{c}
(3) \quad ar(v_1, v_2) \supset nasty(v_1, v_2) \\
(6) \quad \neg \exists x [po(x) \wedge ar(x, p) \wedge fr(x, p) \supset Y] \\
(7) \quad \neg [po(v_1) \wedge ar(v_1, p) \wedge fr(v_1, p) \supset Y] \\
(8) \quad po(v_1) \wedge ar(v_1, p) \wedge fr(v_1, p) \\
(9) \quad \neg Y \\
(10) \quad po(v_1) \\
(11) \quad ar(v_1, p) \\
(12) \quad fr(v_1, p) \\
(13) \quad \neg ar(v_1, v_2) \quad \left| \quad nasty(v_1, v_2) \right. \\
\quad \{v_2 \leftarrow p\} \quad \left| \quad \{Y \leftarrow nasty(v_1, v_2)\} \right. \\
\quad \star \quad \quad \quad \star
\end{array}$$

7 is from 6 by instantiation. Since $\neg(X \supset Y)$ is a conjunctive formula with conjuncts X and $\neg Y$, lines 8 and 9 are added to the tableau (from 7). Similarly, lines 10, 11 and 12 are from line 8 (because line 8 is labelled with a conjunctive formula). At this stage, the disjunctive formula in (3) is used and the tree branches yielding (13). By using the indicated bindings, we derive the contradictions $\neg ar(v_1, p) \wedge ar(v_1, p)$ for the leftmost branch, and $\neg nasty(v_1, p) \wedge nasty(v_1, p)$ for the rightmost one. Since both branches are closed, the tableau is closed and proposition (1) is proved. Thus HOU+R succeeds thereby yielding the appropriate sloppy interpretation for the target clause *So did the one Peter got collared by*.

4 Paycheck Pronouns

Why does the parallelism constraint predict paycheck pronouns? Intuitively, the reason is that as in the ellipsis case, the λ -term shared by source and target utterances may contain the representation of a pronoun. When the antecedent of this pronoun is a parallel element, HOU predicts that this pronoun can behave sloppily. Let us see in more detail how this works. Consider example (27).

Example 27. Jon¹ spent [his₁ paycheck]² but Peter saved it₂.

The pronoun *it* occurring in the second clause has a sloppy interpretation in that it can be interpreted as meaning *Peter's paycheck*, rather than *Jon's paycheck*. In the literature such pronouns are known as *paycheck pronouns* and are treated as introducing a definite whose restriction is pragmatically given (cf. e.g. [2]). Within our proposal, we can straightforwardly capture this intuition by assigning paycheck pronouns the following representation:

$$\text{Pro} \rightsquigarrow \lambda Q. \exists x [P(x) \wedge \forall y [P(y) \rightarrow y = x] \wedge Q(x)] \text{ with } P \in \text{wff}_{(e \rightarrow t)}$$

That is, paycheck pronouns are treated as definites whose restriction (P) is a variable of type $(e \rightarrow t)$. Furthermore, we assume that paycheck pronouns like VP-ellipses, occur in parallel structures and hence are subject to the parallelism constraint. Under these assumption the following equations must hold (we abbreviate $\lambda Q. \exists x [P(x) \wedge \forall y [P(y) \rightarrow y = x] \wedge Q(x)]$ to $\lambda Q. \exists_1 x [P(x) \wedge Q(x)]$):

$$\begin{aligned} \exists_1 x [pc_of(x, j) \wedge sp(j, x)] &= A(j, sp) \\ \exists_1 x [P(x) \wedge sa(p, x)] &= A(p, sa) \end{aligned}$$

Resolving the first equation yields $\lambda y. \lambda O. \exists_1 x [pc_of(x, y) \wedge O(y, x)]$ as a value for A so that $A(p, sa) = \exists_1 x [pc_of(x, p) \wedge sa(p, x)]$ and $\{P \leftarrow \lambda y. pc_of(y, p)\}$. That is, the target clause is correctly assigned the sloppy interpretation: *Peter saved Peter's paycheck*.

5 Source proform: VPE

So far, we have considered cases of sloppy interpretation in which the source proform is a pronoun. Interestingly however, the source proform can also be a VP ellipsis. This is illustrated by the following examples (from [12]).

Example 28. I'll [help you]¹ if you [want me to]². I'll kiss you even if you don't₂.

Example 29. When Harry [drinks]¹, I always conceal [my belief that he shouldn't]₁². When he gambles, I can't conceal it₂.

In (28), the target proform is a VPE whose antecedent contains a VPE. In other words, the source proform is a VPE. Similarly, in (29), the target proform is a pronoun with a VPE as a source proform. Both examples have a sloppy interpretation: *I'll kiss you even if you don't want me to kiss you* (instead of *help you* in the source clause) and *I can't conceal the belief that he shouldn't gamble* (instead of *he shouldn't drink* in the source clause).

What these (and the previous) examples show, is that sloppy interpretation is independent of whether the source and the target proforms are VPEs or pronominal anaphors. A correct theory of sloppy interpretation must therefore be general enough to encompass all possible cases in a uniform way. This is predicted by our account where sloppy interpretation follows, not from the treatment of VPE and paycheck pronouns, but from their interaction with the parallelism constraint. In what follows we show that our approach accounts for examples (28) and (29) thus covering the four possible configurations for sloppy interpretation illustrated in table (1) where boxes surround the source and target proforms. Consider (28) first. Assuming that the parallel elements are *help* and *kiss* respectively, the parallelism constraint requires that

$$\begin{aligned} h(i, you) &\leftarrow wt(you, h(i, you)) = A(h) \\ k(i, you) &\leftarrow P(you) = A(k) \end{aligned}$$

	Source Utterance	Target Utterance
Configuration	$NP^j, \dots, [VP, \dots, \boxed{\text{PRO}}_j, \dots,]^i$	$, \dots, \boxed{\text{VPE}}_i, \dots,$
Example	<i>Jon_j [_{VP} washes his_j car]ⁱ</i>	<i>Peter does_i too.</i>
Configuration	$VP^j, \dots, [VP, \dots, \boxed{\text{VPE}}_j, \dots,]^i$	$, \dots, \boxed{\text{VPE}}_i, \dots,$
Example	<i>I'll [_{VP} help you]^j, if you [_{VP} want me to_j]ⁱ</i>	<i>I'll kiss you even if you don't_i.</i>
Configuration	$NP^j, \dots, [NP, \dots, \boxed{\text{PRO}}_j, \dots,]^i$	$, \dots, \boxed{\text{PRO}}_i, \dots,$
Example	<i>Jon_j spent [_{NP} his_j paycheck]ⁱ</i>	<i>but Peter saved it_i.</i>
Configuration	$VP^j, \dots, [NP, \dots, \boxed{\text{VPE}}_j, \dots,]^i$	$, \dots, \boxed{\text{PRO}}_i, \dots,$
Example	<i>When Harry [_{VP} drinks]_j, I conceal [_{NP} my belief that he shouldn't_j]ⁱ</i>	<i>When he gambles, I can't conceal it_i.</i>

Table 1. Ellipsis, Pronouns and sloppiness

Resolution of the first equation yields $\lambda R.(R(i, you) \leftarrow wt(you, R(i, you)))$ as a possible value for A so that the value for $A(k)$ is:

$$\lambda R.(R(i, you) \leftarrow wt(you, R(i, you)))(k) = A(k)$$

or equivalently $k(i, you) \leftarrow wt(you, k(i, you))$. Indirectly then, the value of P is now $\lambda x.wt(i, k(i, x))$ so that the VPE occurring in the second clause is correctly interpreted as meaning *want me to kiss you*.

Similarly, the derivation for (29) proceeds as follows. This time we assume that *drinks* and *gambles* are the parallel elements so that the parallelism constraint gives rise to the following equations:

$$\begin{aligned} dk(h) &\rightarrow \exists x[bel(x, sn(dk(h))) \wedge hide(i, x)] = A(dk) \\ gb(h) &\rightarrow \exists x[P(x) \wedge hide(i, x)] = A(gb) \end{aligned}$$

Resolving the first equation yields $\lambda R.(R(h) \rightarrow \exists x[bel(x, sn(R(h))) \wedge hide(i, x)])$ as a possible value for A . By β -reduction, the value of $A(gb)$ is then $gb(h) \rightarrow \exists x[bel(x, sn(gb(h))) \wedge hide(i, x)]$ and P is resolved to $\lambda y.bel(y, sn(gb(h)))$. That is the second clause of the target utterance is interpreted as meaning *I can't conceal my belief that Harry shouldn't gamble* and thereby the pronoun *it* is sloppily resolved to *my belief that Harry shouldn't gamble*.

6 Other sloppy constructions

[15] lists a number of constructions where sloppy/strict ambiguity is possible. We now consider these cases. The first construction involves implicit referential arguments and is exemplified in (30).

Example 30. Jon went to a local bar to watch the Super-bowl, and Bob did too.
(*Bob went to a bar local to Bob/Jon*)

If these implicit arguments are taken to be present in the semantic representation, those cases are unproblematic. For instance, the equations for (30) are:

$$\begin{aligned} \exists x[bar(x) \wedge loc(x, j) \wedge went(j, x, watch(j, sb))] &= A(j) \\ R(b) &= A(b) \end{aligned}$$

Solving the first equation yields two solutions for A , namely

$$\begin{aligned} \{A \leftarrow \lambda z.\exists x[bar(x) \wedge loc(x, z) \wedge went(z, x, watch(z, sb))]\} \\ \{A \leftarrow \lambda z.\exists x[bar(x) \wedge loc(x, j) \wedge went(z, x, watch(z, sb))]\} \end{aligned}$$

The first solution yields the sloppy reading and the second, the strict.

Another construction listed in [15] is one involving either pronominal (31) or definite (32) event anaphora.

Example 31. Jon got shot by his father. That happened to Bob too. (*Bob got shot by Bob/Jon's father*)

Example 32. Jon kissed his wife, and Bill followed his example (Dahl, 1972) (*Bill kissed Jon/Bill's wife*)

Here the question arises as to how expressions such as *that happened* and *followed x's example* should be represented. Without going into any details about these constructions, it seems reasonable to assume that these expressions are essentially anaphors in that their meaning is determined by the context. Under this assumption, the second clause of (31) and (32) is then represented as $R(b)$ and the equations for e.g. (31) are:

$$\begin{aligned} \text{shot}(f(j), j) &= A(j) \\ R(b) &= A(b) \end{aligned}$$

As usual, resolution of the first equation gives us two solutions, one sloppy ($\{R \leftarrow \lambda z. \text{shot}(f(z), z)\}$) and one strict ($\{R \leftarrow \lambda z. \text{shot}(f(j), z)\}$).

One-anaphora also permits sloppy/strict ambiguity.

Example 33. Although Jon bought a picture of his son, Bill snapped one himself.

We can capture this by treating one-anaphora similarly to paycheck pronouns that is, as definites whose restriction is pragmatically given. Under this assumption, the equations for (33) are:

$$\begin{aligned} \exists x[\text{pic_of}(x, j's \text{ son}) \wedge \text{bo}(j, x)] &= A(j, \text{bo}) \\ \exists x[P(x) \wedge \text{sna}(b, x)] &= A(b, \text{sna}) \end{aligned}$$

The sloppy interpretation is given by the solution

$$\{A \leftarrow \lambda z \lambda y. \exists x[\text{pic_of}(x, z's \text{ son}) \wedge y(z, x)]\}$$

and the strict reading by the second solution

$$\{A \leftarrow \lambda z \lambda y. \exists x[\text{pic_of}(x, j's \text{ son}) \wedge y(z, x)]\}$$

7 Interaction with syntax and quantification

As should be clear from the above discussion, our analysis predicts a systematic ambiguity between strict and sloppy interpretation. However, there are certain constraints on this ambiguity which result from the interaction of parallelism with other linguistic phenomena. We now consider these constraints.

Syntax Categorical information can affect sloppy/strict ambiguity. Thus although (34) has two readings, one strict and one sloppy, (35) can only have the strict reading.

Example 34. Jon¹'s mother loves him₁ and Bill's mother does too (sloppy/strict)

Example 35. Jon's mother loves Jon and Bill's mother does too (strict only)

Within the standard (that is, non-coloured) HOU framework, such examples are problematic because (34) and (35) have the same semantic representation and therefore, there is no way in which they can be distinguished. In other words, both examples will be predicted to be sloppy/strict ambiguous. However within the HOCU framework, colours can be used to create an interface between semantic construction and other levels of linguistic information. In section 2, we saw that pronouns are variable coloured and hence can give rise either to a strict or to a sloppy interpretation of the target. For full NPs, we stipulate that they be s-coloured. Given this, the equations for (35) are:

$$\begin{aligned} l(m(j_p), j_s) &= A_B(j_p) \\ R_s(m(b_p)) &= A_B(b_p) \end{aligned}$$

Crucially, resolution of the first equation yields only one solution, not two, namely $A_B = \lambda z.l(m(z), j_s)$ – this yields the strict reading. By contrast, the sloppy reading is ruled out because the corresponding substitution $\{A_B \leftarrow \lambda z.l(m(z), z)\}$ is not a unifier for the given equations. To see this, it suffices to apply this substitution to the right-hand side of the equation, $A_B(j_p)$. This yields $\lambda z.l(m(z), z)(j_p)$ which by β -reduction is equivalent to $l(m(j_p), j_p)$. But this does not unify with the left-hand side of the equation because of the colour-clash on the second occurrence of j . Hence the sloppy reading is ruled out.

Scope constraints Quantification also constrains sloppiness. This is illustrated by the following examples.

Example 36. Every¹ man believes he₁ is a fool. No, PETER believes he is a fool. (sloppy only)

Example 37. Jon lost a¹ book and never got it₁ back. No, Peter lost a PEN and never got it back (sloppy only)

In both (36) and (37), the source pronoun has a quantified antecedent and furthermore, this quantified NP does not scope over the target utterance. As a result, only the sloppy interpretation is possible. The strict interpretation is ruled out because it would involve a free variable, the variable introduced by the source quantified antecedent and occurring in the target representation. Under our analysis, this constraint simply follows from general constraints on substitutions which essentially say that a free variable may never become bound and vice-versa. Thus for (36), the parallelism constraint requires that the following equations hold (as in DSP, when a proper name has a quantified parallel element, we use the type-raised representation of proper names to guarantee parallelism of types):

$$\begin{aligned} \forall x[b(x) \rightarrow bel(x, f(x))] &= A(\lambda P.\forall x[b(x) \rightarrow P(x)]) \\ bel(p, f(y)) &= A(\lambda P.P(p)) \end{aligned}$$

To obtain the strict reading, the second equation needs to be resolved as follows:

$$\{A \leftarrow \lambda Q.Q(\lambda z.bel(z, f(x)), y \leftarrow x)\}$$

Applying this substitution to the first equation, we then get

$$\begin{aligned} A(\lambda P.\forall x[b(x) \rightarrow P(x)]) &= \lambda Q.Q(\lambda z.bel(z, f(x))(\lambda P.\forall x[b(x) \rightarrow P(x)]) \\ &= \lambda P.\forall x[b(x) \rightarrow P(x)](\lambda z.bel(z, f(x))) \end{aligned}$$

And at this point, the restrictions on well-formed substitutions requires that x be renamed. Hence x can never be bound by the source quantifier *every man*.

Binding constraints As is illustrated by the following example, the binding constraints also play a role in constraining strict interpretation.

Example 38. Mary persuaded Jon¹ to shave himself₁. No, Mary persuaded PETER to shave himself. (sloppy only)

Here only the sloppy interpretation is possible. The most obvious explanation for the missing strict reading (*Peter shaved Jon*) is that the resulting structure would violate condition A of the binding theory according to which a reflexive must be bound in its governing category. Now recall that we only require *discourse* pronouns to be represented by free variables. In contrast, we take reflexive and bound pronouns to be resolved at the sentential level and hence represented either as constants or as bound variables. Given this, the lack of strict reading in (38) above, straightforwardly follows from the parallelism constraint. The equations are:

$$\begin{aligned} pe(m, sh(j, j)) &= A(j) \\ pe(m, sh(p, p)) &= A(p) \end{aligned}$$

So that the only possible solution is $\{A \leftarrow \lambda x.pe(m, sh(x, x))\}$. In particular, the solution which would yield a strict reading namely, $\{A \leftarrow \lambda x.pe(m, sh(x, j))\}$ is not a unifier since it cannot solve both equations simultaneously.

8 Comparison with other approaches

In this section, we briefly compare it with three alternative proposals: DSP's treatment of ellipsis (because it is closely related to our proposal); Hardt's dynamic theory of ellipsis (because it uses a dynamic rather than a static semantics and is in this sense, fundamentally different from our proposal) and Fiengo and May's LF approach (because contrary to our analysis which works on flat semantic representations, this approach involves structured objects namely, logical forms).

Let us first examine how our proposal relates to DSP's treatment of ellipsis. In DSP's analysis, the semantic representation of the source clause is constrained to be equal to $R(S_1, \dots, S_n)$ where R is the semantic representation of the target

ellipsis and S_1, \dots, S_n are the semantic representations of the source parallel elements (that is, the elements of the source clause which have an overt parallel counterpart in the target clause). Reformulating DSP’s analysis in a way that makes it more comparable to our proposal, we then have that the interpretation of an elliptical clause must obey the following constraint:

$$\begin{aligned} S &= R(S_1), \dots, (S_n) \\ T &= R(T_1), \dots, (T_n) \end{aligned}$$

where R is the semantic representation of the target ellipsis; S and T represent the source and the target clause and $S_1, \dots, S_n, T_1, \dots, T_n$ represent the source and the target parallel elements respectively.

This looks a lot like our parallelism constraint. There is one important difference though: DSP’s analysis requires that the semantically underspecified element whose value is to be determined by HOU, be the semantic representation of a VP-ellipsis. By contrast, we only require that this semantically underspecified element be the semantics shared by two parallel propositions.

This difference has several important consequences. First, it allows the treatment of deaccenting (because the free variable R may represent overt rather than elided material). Second, it captures the extended domain of licensing involved in both ellipsis and deaccenting context (because the free variable need not represent an ellipsis but may extend over overt material *containing* an ellipsis). Third, it enables a general theory of sloppy identity (since sloppiness is linked to parallelism, not just ellipsis). Fourth, it allows us to preserve the assumption that VPEs denote properties (in DSP’s analysis, the variable representing the VPE is a relation of varying arity). This allows compatibility with Montague-type grammars that is, grammars where the semantic type of a constituent is defined by a mapping from syntactic categories into types.

We now turn to Hardt’s approach. In a dynamic setting, anaphors can be viewed as denoting functions from contexts to semantic objects. Specifically, an ellipsis can be viewed as a function from contexts to properties. As [7,12] show, this simple observation suffices to predict sloppy/strict ambiguity: if the antecedent of an ellipsis contains a pronoun, the value of this pronoun in the source context may differ from its value in the target context. [12] further shows that in a dynamic setting where pronouns and ellipsis are uniformly treated as function from contexts to semantic objects, the free interplay between pronouns, ellipsis and sloppy identity discussed in section 5 simply falls out.

There are at least two main differences between Hardt’s and our proposal. First, Hardt’s proposal does not impose a parallelism constraint on sloppy identity. This has pros and cons. On the positive side, this means that the approach encounters no difficulty in accounting for examples such as (22), where the non-contrastive material in the target differs from its parallel counterpart in the source. However, it also means that sloppy identity is unrestricted. In particular, a sloppy interpretation for (12) cannot be ruled out.

The second difference involves directionality. Dynamic semantics is inherently directional in that within the semantic representation, the term representing an anaphor must be preceded by the term representing its antecedent. As a result,

the dynamic approach correctly predicts that (39) lacks a sloppy interpretation that is, the reading *this year, I voted for Harry* (example from [12]). Note that by contrast, the HOU approach will predict both a strict and a sloppy reading for the elliptic clause.

Example 39. Tom¹ is always causing me problems. Last year, I didn't vote for him₁ and Tom got mad at me. This year, I did, and HARRY got mad at me.

The blessing is mixed however, for consider the following example.

Example 40. Tom is always causing me problems. Last year, he got mad at me because I didn't vote for him and THIS year, HARRY got mad at me because I did.

In this case, *Harry* does precede the ellipsis and consequently, the dynamic approach predicts a sloppy reading. However, just as in (39), this reading is simply not there. In short, I don't think example (39) is a very decisive argument in favour of the dynamic approach. As example (40) shows, the lack of sloppy reading is due not to linear order, but to some other factor, possibly the semantics of the contrast relation which holds between target and source or possibly, the fact that *Tom* is the current discourse topic.

Finally, we consider [5]'s account (henceforth, FM). Under this account, ellipsis is taken to involve syntactic reconstruction and furthermore, different types of NPs are represented differently. For instance, a proper name is always represented by an α -occurrence but a pronoun can be represented either as a β - or as an α -occurrence. Importantly, reconstruction is sensitive to the $\alpha - \beta$ distinction so that specific constraints on reconstruction can be stated which in essence, aim at capturing the effect of binding on ellipsis reconstruction. In particular, FM's approach predicts the difference between full NPs and pronouns discussed in section 7. It also accounts for the fact that example (41) lacks a strict-sloppy reading that is, the reading *Peter said Max saw Peter's mother*.

Example 41. Max¹ said he₁ saw his₁ mother and Peter² did too

The HOCU approach presented here cannot explain this missing reading for the simple reason that Co-indexed pronouns have identical representations and are thus handled identically. Although I do not think the problem unsolvable, I will leave it for now as an open issue.

9 Conclusion

We have presented a uniform treatment of sloppy interpretation which is general enough to cover the cases of sloppy/strict ambiguity observed in the literature. Moreover the treatment integrates a parallelism constraint which restricts sloppy interpretation to those cases where the antecedent of the source proform has a parallel counterpart in the target clause.

There are two obvious directions for further research. The first direction concerns the determination of parallelism. A real weakness of the present paper

is that parallel elements are taken as given. This undermines the predictive power of the approach in that there is a definite leeway in deciding what the parallel elements can or may be. Roughly, there are two types of approach to this problem. Either, the parallel elements are determined through some general constraint on parallelism and contrast (cf. e.g. [13]), or they are defined through some matching mechanism on semantic representations such as priority union (cf. e.g. [17,11]). Both have their merits but it seems fair to say that while the general discourse level approach is too vague to be useful for our purpose, the more precise approach advocated by the matching proposals is too restrictive to be empirically very appropriate. To attain a theory of parallelism structuring, an integration of both methods is called for, which combines the general insights of discourse theories of parallelism with the precision of a given matching operation.

The second line of research which needs pursuing concerns the interaction of sloppy/strict ambiguity with other linguistic phenomena. In section 7, we have seen that colours allow for the integration of non-semantic information into the semantic construction process. However, we have also seen cases in which colours are not enough (cf. section 8). In particular, our analysis fails to predict the ‘eliminative puzzles of ellipsis’. That is, it fails to capture the fact that in discourses with n referring elements in the source, less than 2^n readings are actually available (cf. [3,5]) and some readings are ruled out. For these cases and more generally, for a full treatment of how semantic construction interacts with other levels of linguistic information, a more sophisticated apparatus is needed. This is the subject of another paper.

Acknowledgments

I would like to thank Daniel Hardt for comments and criticisms on previous versions of this paper. The work reported here was funded by the Deutsche Forschungsgemeinschaft in Sonderforschungsbereich SFB–378, Project C2 (LISA).

References

1. Noam Chomsky and Howard Lasnik. *Syntax: An International Handbook of Contemporary Research*, chapter Principles and Parameter Theory. Walter de Gruyter, 1991. To appear.
2. Robin Cooper. The interpretation of pronouns. In F. Heny and H.S. Schnelle, editors, *Syntax and Semantics*, number 10, pages 61–93. 1979.
3. Östen Dahl. On so-called sloppy identity. Gothenburg Papers in Theoretical Linguistics 11, University of Göteborg, 1972.
4. Mary Dalrymple, Stuart Shieber, and Fernando Pereira. Ellipsis and higher-order unification. *Linguistics and Philosophy*, 14:399–452, 1991.
5. Robert Fiengo and Robert May. *Indices and Identity*. MIT Press, Cambridge, 1994.
6. Melvin Fitting. *First-Order Logic and Automated Theorem Proving*. Springer Verlag, 1990.
7. Claire Gardent. Dynamic semantics and vp ellipsis. In Jan van Eijck, editor, *Logics in AI*, pages 251–267, Amsterdam, The Netherlands, 1990. Springer-Verlag.

8. Claire Gardent. Parallelism, *hou* and deaccenting. CLAUS Report 85, University of Saarbrücken, January 1997.
9. Claire Gardent and Michael Kohlhase. Higher-order coloured unification and natural language semantics. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*. ACL, Santa Cruz, 1996.
10. Claire Gardent, Michael Kohlhase, and Noor van Leusen. Corrections and Higher-Order Unification. In *Proceedings of KONVENS96*, pages 268–279. De Gruyter, Bielefeld, Germany, 1996.
11. Claire Grover, Chris Brew, Suresh Manandhar, and Marc Moens. Priority union and generalisation in discourse grammar. In *Proceedings of the 31th Annual Meeting of the Association for Computational Linguistics*, 1994.
12. Daniel Hardt. Dynamic interpretation of verb phrase ellipsis. Ms., University of Villanova, 1996.
13. Jerry Hobbs. *Literature and Cognition*. Number 21 in CSLI Lecture Notes. CSLI, 1990.
14. Dieter Hutter and Michael Kohlhase. A coloured version of the λ -calculus. In William McCune, editor, *Proceedings of the 14th Conference on Automated Deduction*, LNAI, Townsville, Australia, 1997. Springer Verlag. forthcoming.
15. Andrew Kehler. *Interpreting Cohesive Forms in the Context of Discourse Inference*. PhD thesis, Harvard University, 1995.
16. Michael Kohlhase. Higher-order tableaux. In *Proceedings of the Tableau Workshop*, Koblenz, Germany, 1995.
17. H. Prüst, R. Scha, and M. van den Berg. Discourse grammar and verb phrase anaphora. *Linguistics and Philosophy*, 17:261–327, 1994.
18. Mats Rooth. Ellipsis redundancy and reduction redundancy. In Steve Berman and Arild Hestvik, editors, *Proceedings of the Stuttgart Ellipsis Workshop*, University of Stuttgart, 1992.
19. Christopher Tancredi. *Deletion, Deaccenting and Presupposition*. PhD thesis, Massachusetts Institute of Technology, 1992.